arXiv:1609.01558v2 [quant-ph] 7 Feb 2017

Towards optimal experimental tests on the reality of the quantum state

George C. Knee

Department of Materials, University of Oxford, Oxford OX1 3PH, United Kingdom
Department of Physics, University of Warwick, Coventry CV4 7AL, United Kingdom

E-mail: gk@physics.org

8 February 2017

Abstract. The Barrett-Cavalcanti-Lal-Maroney (BCLM) argument stands as the most effective means of demonstrating the reality of the quantum state. Its advantages include being derived from very few assumptions, and a robustness to experimental error. Finding the best way to implement the argument experimentally is an open problem, however, and involves cleverly choosing sets of states and measurements. I show that techniques from convex optimization theory can be leveraged to numerically search for these sets, which then form a recipe for experiments that allow for the strongest statements about the ontology of the wavefunction to be made. The optimization approach presented is versatile, efficient and can take account of the finite errors present in any real experiment. I find significantly improved low-cardinality sets which are guaranteed partially-optimal for a BCLM test in low Hilbert space dimension. I further show that mixed states can be more optimal than pure states.

1. Introduction

The ontological status of the quantum state has long been a central question in foundational physics. If it is *ontic*, and therefore a true part of physical reality, then the many counterintuitive features of quantum theory (QT) remain opaque. The idea that it is *epistemic* – that is, reducible in essence to a state of knowledge – is an attractive proposition that promises to dissolve some of these issues [1]. For instance, the instantaneous and discontinuous 'collapse' of the wavefunction is arguably much more naturally thought of as a Bayesian update of knowledge than as a process governed by dynamical physical laws. There are a plethora of other canonically 'quantum' phenomena that have an appealing explanation when adopting an epistemic interpretation – see for example Ref. [2]. In recent years, there has been a flurry of progress towards understanding the feasibility of the so called 'ψ-epistemic' programme [3], culminating in experimental tests [4]. Here I improve on the design of such experiments, so that tighter restrictions on possible ψ-epistemic theories may be determined in the laboratory.

In the ontological-models framework, where these notions are made precise, the preparation of a quantum state $|\phi\rangle$ is associated with the (generally random) selection of an ontic state $\lambda \in \Lambda$, with the appropriately normalised distribution of probability for the various states being written $\mu_\phi(\lambda)$. Λ is the space of ontic states, and I denote the subset Λ_ϕ as the support of μ_ϕ – members of which are said to be 'compatible' with the preparation. A projective measurement operator $|\psi\rangle\langle\psi|$, on the other hand, is associated with a conditional probability $\xi(\psi|\lambda)$ known as a response function. In the full framework, there are also stochastic maps on Λ which represent, e.g. unitary transformations, but it will not be necessary to consider these here.

In order to constitute a model of quantum theory, the following condition must be met

$$\int \mu_\phi(\lambda)\xi(\psi|\lambda)d\lambda = |\langle\psi|\phi\rangle|^2. \tag{1}$$

The ontic state λ, therefore, stands for a variable sufficient to screen off the preparation μ from the measurement ξ. As long as (1) is satisfied, tracing over the ontic states leaves the 'correct' conditional probabilities *à la* the Born rule. Many such ontological models are possible, and they may be classified according to certain properties of the μ and ξ. Although couched in the older terminology of 'hidden-variables', one of the most important such classifications is Bell's definition of locality [5], from which he derived an experimentally testable inequality able to separate certain predictions of quantum theory from all possible *local* ontological-models [6]. Bell was therefore perhaps the first to show that a no-go theorem can provide positive insight into what an ontological model can be. Importantly, Bell's inequality was made robust to experimental error [7], and has now been subject to very strict tests [8, 9, 10]. Since QT could well be false, theoretical proofs which rely on quantum predictions do not establish (for example) the non-locality of nature itself. Experimental tests are therefore highly important: since they can indeed establish such features in a theory independent manner.

More recent work divides the ontological models into those where the wavefunction is *ontic* $\rightarrow \Lambda_\psi \cap \Lambda_\phi = \emptyset \quad \forall\psi \neq \phi$, from those which are *not-ontic*, or *epistemic* $\exists\psi \neq \phi : \Lambda_\psi \cap \Lambda_\phi \neq \emptyset$. Pusey, Barrett and Rudolph (PBR) proved a theorem that theoretically ruled out the subset of epistemic models where the preparation of two independent systems can be assumed to be represented by the product $\mu_1 \times \mu_2$ [11]. There are two reasons why the epistemic view cannot be immediately dispensed with, however. First, even in the idealized case, one may always instead drop the assumption of preparation independence (PI). Indeed, constructive examples of ψ-epistemic theories exist [12, 13, 14] which must therefore dodge PBR's no-go theorem in this way. Second, the finite precision of any real experiment, including a recent one performed in an ion trap [15], mean that only a strict subclass of epistemic models subscribing to PI are ever experimentally falsified – the 'wriggle room' offered by laboratory imperfections makes it possible to retain *both* PI *and* an epistemic view of nature. Subsequent theoretical studies have argued for the reality of the quantum state by imposing further assumptions on the set of ontological models [16, 17, 18] – for a comprehensive review, see Ref. [3].

Maroney introduced a classification of ontological models which generalises the ontic/epistemic dichotomy [19], so that one can begin to constrain the extent to which a model may be epistemic. Crucially, Maroney's theorem does not rely on PI, and thus its implications cannot be escaped by discarding it. Maroney's idea, later made noise tolerant by Barrett et al (BCLM) [20], springs from a very particular motivation for ψ-epistemicism: namely, the impossibility of discriminating non-orthogonal quantum states. In quantum theory, this feature depends on the quantity

$$\omega_Q(|\psi\rangle, |\phi\rangle) := 1 - \sqrt{1 - |\langle\psi|\phi\rangle|^2}; \tag{2}$$

a measure of the overlap of the quantum states in Hilbert space. It is related to the maximum probability with which the two states can be distinguished in a single shot experiment [21]. An analogous quantity

$$\omega_C(\mu_\psi(\lambda), \mu_\phi(\lambda)) := 1 - \frac{1}{2}\int |\mu_\psi(\lambda) - \mu_\phi(\lambda)| d\lambda \tag{3}$$

applies to the underlying ontological model: it measures the extent to which two preparations 'overlap' in the ontic statespace. A partially-ψ-epistemic model is defined by the relation:

$$k(\psi, \phi) := \frac{\omega_C(\mu_\psi(\lambda), \mu_\phi(\lambda))}{\omega_Q(|\psi\rangle, |\phi\rangle)}, \quad \langle\psi|\phi\rangle \neq 0. \tag{4}$$

The proportionality constant $k(\psi, \phi)$ will be the main subject of our study. If $k(\psi, \phi) = 0$, there is no overlap of distributions in the ontic statespace, irrespective of the 'closeness' of the quantum states they represent: the model is ψ-ontic. On the other hand, since it is required that $\omega_C(\mu_\psi(\lambda), \mu_\phi(\lambda)) \leq \omega_Q(|\psi\rangle, |\phi\rangle)$ [15], a model with $k(\psi, \phi) = 1$, is said to be 'maximally ψ-epistemic'.

The remainder of this paper is structured as follows: in Section 2, I provide an operational interpretation for intermediate values of $k(\psi, \phi)$; in Sections 3 and 4 I recapitulate the BCLM argument and describe the search for optimal experimental implementations as a nonlinear optimization problem. I present an algorithm for this purpose in Section 4.1, along with several numerical results in Section 5. There I show how the algorithm can tailor the design of an experiment by accepting the typical error rate as input (in Section 5.1). Conclusions are drawn at the end of the paper, and the appendices contain details of a) an algorithmic subroutine, b) a mixed-state result that performs more optimally than any known set of pure states of the same size, and c) algorithm runtime statistics.

2. Degrees of epistemicness

The task of discriminating $|\phi\rangle$ from $|\psi\rangle$ in a single shot has a minimum probability of making an error (i.e. guessing the wrong state) given by $P_Q^{\text{error}} = \frac{\omega_Q}{2}$; relating to the case where one makes use of the best measurement available under the constraints set by quantum theory [22]‡. A hypothetical *omniscient* measurement is described by

‡ When referring to minimum error distinction probability, we assume throughout that there is no prior information available to the state discriminator.

a continuous set of Dirac delta functions $\xi(\lambda'|\lambda) = \delta(\lambda' - \lambda)$: it strongly violates the aforementioned constraints and reveals complete and perfect information about λ. Such a measurement, which is the 'best' that is logically possible, has an error probability

$$P_{\infty}^{\text{error}} = \frac{\omega_C}{2} = k(\psi,\phi)\frac{\omega_Q}{2} = k(\psi,\phi)P_Q^{\text{error}}. \tag{5}$$

$k(\psi,\phi)$ therefore represents the improved state-discrimination error-probability of an omniscient observer, compared with the observer whom is constrained by quantum mechanics. In situations where the former would be superior, this necessitates a certain property of the response functions in the ontological model termed 'deficiency' by Harrigan and Rudolph [23]. Deficiency means a measurement of $|\psi\rangle\langle\psi|$ will respond to ontic states λ which are not compatible with μ_ψ. If $k(\psi,\phi) = 0$, an ontological model of quantum theory must be maximally deficient – that is to say, the inequality

$$\int_{\Lambda\setminus\Lambda_\psi} \mu_\phi(\lambda)\xi(\psi|\lambda)d\lambda \leq \int_{\Lambda} \mu_\phi(\lambda)\xi(\psi|\lambda)d\lambda = |\langle\psi|\phi\rangle|^2 \tag{6}$$

will be saturated. A fuller discussion of deficiency is given in [24].

3. Bounds on $k(\psi,\phi)$

BCLM showed how $k(\psi,\phi)$ can be bounded without relying on PI [20]. In Hilbert space dimension d, consider a reference state $|c\rangle$ along with a set of n states $\mathcal{S} = \{|\psi_i\rangle\}_{i=1}^n$. Consider further a set of projective operators $\mathcal{M} = \{|ijk\rangle\langle ijk|; i = 1\ldots n, j = 1\ldots i-1, k = 1,2,3\}$, such that associated to each pair of states $\{|\psi_i\rangle, |\psi_j\rangle\}$ in $\mathcal{S}$, there is a triple of orthogonal projectors $\{|ij1\rangle\langle ij1|, |ij2\rangle\langle ij2|, |ij3\rangle\langle ij3|\}$ in $\mathcal{M}$ that define a three-outcome measurement. We will find it useful to define $\mathcal{A}_{ij} := |\langle ij1|\psi_i\rangle|^2 + |\langle ij2|\psi_j\rangle|^2 + |\langle ij3|c\rangle|^2$, which is a measure of the *antidistinguishability* of the triple of states $|\psi_i\rangle, |\psi_j\rangle, |c\rangle$. This concept is at the heart of many ψ-ontology theorems: when $\mathcal{A}_{ij} = 0$ the triple is said to be perfectly antidistinguishable, and then the measurement $\{|ijk\rangle\langle ijk|\}_{k=1,2,3}$ will conclusively exclude (in a single shot) the possibility that one of the triple of states was prepared [25]. Caves, Fuchs and Schack have derived the necessary and sufficient conditions for perfect antidistinguishability (which they call PP-incompatibility) of a triple of states, which depend only on the three inner products between pairs in the triple [26].

Next, define $k_0 := \min_j k(c,\psi_j)$. Then, if QT predictions are correct, the relation

$$k_0 \leq \frac{1 + \sum_{i>j} \mathcal{A}_{ij}}{\sum_i \omega_Q(|c\rangle, |\psi_i\rangle)} \tag{7}$$

follows from (4) and from the Bonferonni inequality, as BCLM show [20]. The remainder of this paper is concerned with finding small $\mathcal{S}$ and $\mathcal{M}$ such that if QT is approximately correct, they would lead to the lowest upper bound on k_0. An upper bound $k_0 \leq x$ implies that there exists at least one state in $\mathcal{S}$ which has a degree of epistemicness with respect to c no greater than x. One may optionally make additional assumptions to extrapolate this to a stronger claim: for example, Lipschitz continuity [3] would require

that $k(\psi, \phi) \leq x \quad \forall \psi, \phi$. In the course of our search, we will limit ourselves to finite values for d the Hilbert space dimension and n the number of states, considered here as experimental resources to be spent frugally.

3.1. Existing families of states

Before presenting my results below, it is prudent to survey currently-available solutions the problem which will serve as benchmarks for our algorithmic approach. BCLM supplied for every $d \geq 4$ and also power prime, a set of $n = d^2$ states, satisfying $\mathcal{A}_{ij} = 0 \quad \forall i, j$ which leads to the bound $k_0 < 2/d$ [20]. When $d \geq 4$ is not power prime, their results lead to the bound $k_0 < 4/(d-1)$. Relaxing the need for exact antidistinguishability enabled them to find $k_0 \leq 0.95$ for $d = 3, n = 9$. Next, Leifer showed an exponential decay in d (for $d \geq 3$) by using a set of $n = 2^{d-1}$ Hadamard states [27]: this achieves the bound $k_0 < 4d/2^d$ [28]. Branciard provided a nonconstructive proof that $k_0 \leq 8/n^{(d-3)/(d-2)}$, which displays a decay in n for any $d \geq 4$, as well as a number of constructive solutions [28]. Where these previous results provide bounds on k_0 for low n and d, the approach I present below was able to match or better the bound. The best bound achieved experimentally at the time of writing is $k_0 \lessapprox 69\%$ [4]. Achieving $k_0 \leq 50\%$ would seem like the next major milestone, where the classical overlap is doing less than half of the necessary work in explaining the indistinguishability of non-orthogonal quantum states.

4. Optimization

We wish to solve

$$\begin{aligned} \underset{\mathcal{S}\times\mathcal{M}}{\text{minimize}} \quad & \frac{1 + \sum_{i>j} \mathcal{A}_{ij}}{\sum_i \omega_Q(|c\rangle, |\psi_i\rangle)} \\ \text{subject to} \quad & \langle \psi_i | \psi_i \rangle = 1 \quad i = 1, \ldots n \\ & \langle ijk | ijk \rangle = 1 \quad i = 1, \ldots, n \quad j = 1, \ldots, i \quad k = 1, 2, 3 \\ & \langle ij1 | ij2 \rangle = \langle ij1 | ij3 \rangle = \langle ij2 | ij3 \rangle = 0 \quad i = 1, \ldots, n \quad j = 1, \ldots, i \end{aligned} \tag{8}$$

Here all optimisation variables are considered as unnormalised vectors in $\mathbb{C}^d$, with the necessary normalisation and orthogonality conditions shown explicitly above as constraints. An analytic, global solution to this problem seems intractable: but how difficult is the numerical optimization problem at hand? A naive answer is found by counting the number of real parameters necessary to describe a solution to the problem. One can do better than simply taking the real and imaginary components of each vector, since some of the constraints imply these parameters are not independent. For example, the hyper-spherical parameterisation of pure states requires only $d-1$ polar angles and $d-1$ phase angles [29] (making $2d-2$ in all), with normalisation constraints automatically satisfied. A projective measurement is defined by a $d \times d$ unitary matrix: in close analogy to the argument above concerning states, a unitary

matrix may be parameterised by d^2-1 angles [29], and describes an orthonormal measurement basis by construction. Thus, even with a smart parameterisation one must solve a $(2d-1)n+(d^2-1)(n^2-n)/2$ dimensional optimization problem.

Various heuristics can inform our search. For example, to make each $|\psi_i\rangle$ as close as possible to $|c\rangle$, whilst at the same time making each state far from all of the others so that they can be (approximately) anti-distinguished. The objective function is nonlinear, nonconvex and is pocked with local optima: gradient descent will therefore most likely get 'stuck' in a feasible region with sub-par performance. Brute force methods would be intractable: note that, on a coarse grid dividing each angle into only g discrete values, one must evaluate the objective function g^{36} times in the problem instance of $d=n=3$. Global approaches include simulated annealing, particle swarm and associated techniques. But a more powerful approach can be implemented by capitalising on a specific structure of the problem.

4.1. Exploiting convexity

One may 'lift' the problem from a search over the set of n statevectors $\mathcal{S}$ to the set of n density matrices $\mathcal{D}=\{\rho_i\}_i$ and from the set of $3(n^2-n)/2$ projective operators $\mathcal{M}$ to the set of the same number of positive operators $\mathcal{P}=\{E_{ijk}\}_{ijk}$. Let us take a moment to interpret what this might mean. A Positive Operator Valued Measure (or POVM) is a set of positive operators that sum to the identity operator. It constitutes the most general description of a quantum measurement. Likewise, a mixed state (being an Hermitian, trace-one operator) is the most general description of a quantum state. If we interpret a mixed state as a proper mixture of other states, then the corresponding preparation in the ontological model is the random selection of one of a number of preparations, each itself a random selection of λ. See Figure 1.

The following generalisations of the quantum overlap and the Born rule

$$\omega_Q(|c\rangle,|\psi_i\rangle) \to \omega_Q(\rho_c,\rho_{\psi_i}) = 1-\frac{1}{2}||\rho_c-\rho_i||_* \tag{9}$$

$$|\langle ijk|\psi_c\rangle|^2 \quad \to \mathrm{trace}(E_{ijk}\rho_c) \tag{10}$$

share an important property: *they conserve their meanings* as twice the minimum error discrimination probability [21], and prepare-measure probability (respectively).$||\bullet||_*$ is the nuclear norm (or sum of singular values). Derivative notions such as antidistinguishability $\mathcal{A}_{ij}$ therefore also inherit their meaning correspondingly. With the generalisations in place, we can therefore write the optimization problem as

$$\begin{aligned}
\underset{\mathcal{D}\times\mathcal{P}}{\text{minimize}} \quad & \frac{1+\sum_{i>j}\mathrm{trace}(E_{ij1}\rho_i+E_{ij2}\rho_j+E_{ij3}\rho_c+\epsilon_{ij1}+\epsilon_{ij2}+\epsilon_{ij3})}{\sum_j 1-\frac{1}{2}||\rho_c-\rho_j||_*} \\
\text{subject to} \quad & E_{ij1}+E_{ij2}+E_{ij3}=\mathbb{I} \\
& E_{ijk}\succeq 0 \\
& \rho_i\succeq 0 \\
& \mathrm{trace}(\rho_i)=1.
\end{aligned} \tag{11}$$

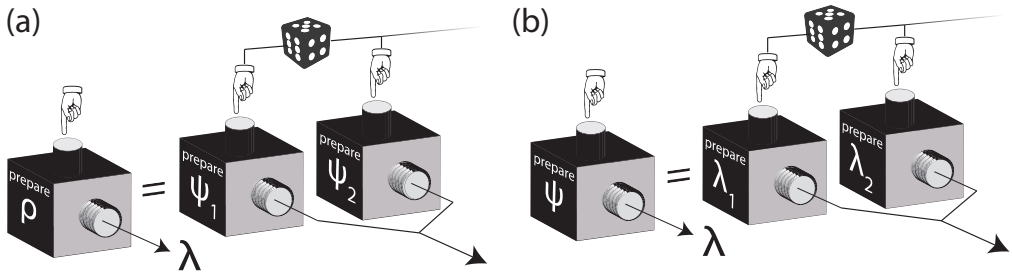


Figure 1. In the ontological models framework, preparations may be thought of as 'black boxes' that simply produce a λ according to the distribution $\mu_\psi(\lambda)$. (a) A mixed state $\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|$ can be thought of as the net preparation when several pure preparations are wired together with a probabilistic switch. The resulting probability density over ontic states is the convex combination of the component distributions: $\mu_\rho(\lambda) = \sum_i p_i \mu_{\psi_i}(\lambda)$. (b) In much the same way, each component preparation labelled by a pure state ψ_i can be thought of itself as a compound preparation, where a probabilistic switch selects from a number of deterministic preparations of λ, so that the net preparation is described by $\mu_\psi(\lambda)$. For the purposes of illustration, here Λ is represented as a discrete space; although it is often (and in eq (1)) thought of as a continuous space.

Here $M \succeq 0$ denotes that M is a positive semidefinite matrix. ϵ_{ijk} are small error parameters to be discussed shortly – the reader may temporarily assume these to be zero. Recall that a function $f(\mathbf{x})$ is convex iff $f(\lambda\mathbf{x}_1+(1-\lambda)\mathbf{x}_2) \leq \lambda f(\mathbf{x}_1)+(1-\lambda)f(\mathbf{x}_2)$, and that a set $\mathcal{B}$ is convex iff $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{B} \rightarrow \lambda\mathbf{x}_1 + (1-\lambda)\mathbf{x}_2 \in \mathcal{B}$ for $\lambda \in [0,1]$. Call a problem *convex* if it involves minimising a convex function over a convex set. Such problems exhibit many convenient features: if a local optimum exists, it is also a global optimum. Efficient algorithms exist to solve convex problems. Although (11) is not a convex problem, notice that the objective function is linear (hence convex) in the set of measurements when the states are fixed. Furthemore, when the measurements are fixed, the objective is convex-concave fractional in the states: then the global optimum can be found by solving a series of parametric (convex) subproblems [30]. Note further that $\mathcal{D}, \mathcal{P}$ are convex sets. Thus, we proceed in a manner inspired by 'biconvex' problems [31] (which have a very similar structure). The approach put forward here is to begin with a feasible point in $\mathcal{D}$, and then proceed to alternately optimise over $\mathcal{P}$ and then $\mathcal{D}$ again, and so on (keeping the other set fixed) until the optimal values of the two subproblems converge. This is known as alternate convex search (ACS) [31], and whilst it does not guarantee global optimality, it does tend to provide good results which are guaranteed *partially optimal*: that is to say, no change in $\mathcal{D}$ or in $\mathcal{P}$ alone could provide a better solution. In order to solve the fractional subproblem of the form min a/b (when fixing the measurements and searching for states), one may use Dinkelbach's technique [32], which involves solving a series of parameterised problems min $a - \theta_k b$: see Appendix A.

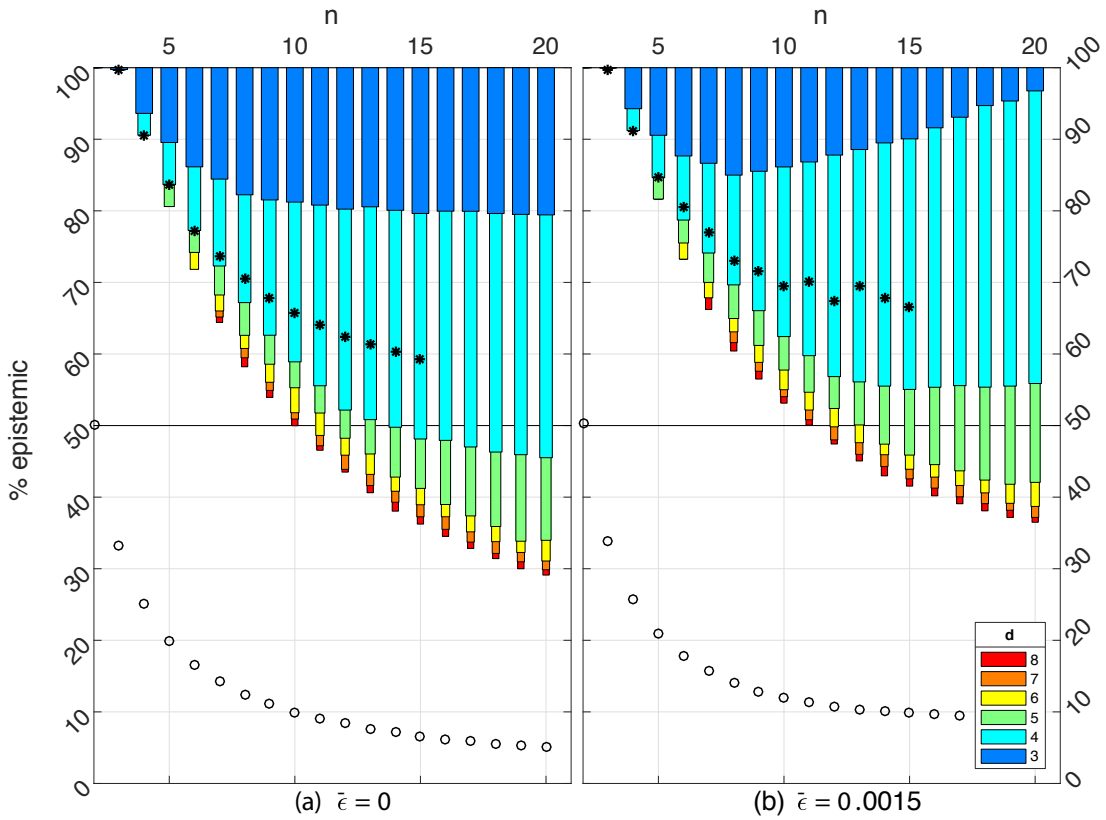


Figure 2. (a) Upper bounds on k_0 achievable by preparing $n+1$ quantum states each of dimension d, generated using an alternating search which solves convex subproblems. Increasing d from 3 to 4 and to 5 provide significant gains, but thereafter only diminishing returns. Open circles follow $(1+\frac{1}{2}n(n-1)3\bar{\epsilon})/n$, which is a lower bound to the optimum experiment (independent of d). (b) Allowing observed statistics to deviate on average by 0.15% from ideal QT predictions defines a new optimization problem, with more modest bounds still achievable. Note how an optimum n emerges. Black asterisks indicate previously best-known a) theoretical and b) experimental bounds from Ringbauer et al. in $d = 4$ (with $\bar{\epsilon} \sim 0.0015$) [4].

5. Numerical results

There exist many software packages for solving the convex problems that arise during my algorithm. Because the algorithm will not always settle on the same solution if seeded with different starting points, I ran the algorithm several times, using the CVX package [33] for each subproblem: the best numerical results for $\epsilon_{ijk} = 0$, $d = 3, \ldots, 8$ and $n = 3, \ldots, 20$ are shown in Figure 2a. The algorithm was seeded with random pure states. Corresponding MATLAB files with the states and measurements are available in the supplementary material, as is the code needed to find partially optimal solutions for any n, d. The results show that an experiment showing $k_0 \leq 50\%$ is possible with current technology such as linear optics [4] or ion traps [15]. The performance of the algorithm is discussed in Appendix C.

Some comments about the choice of reference state are in order. Because our

objective function is invariant under a simultaneous unitary transformation of all states and measurements, the reference state may be chosen arbitrarily *up to the choice of eigenvalues*, which determine (for example) the purity trace(ρ_c^2). Because of the linearity of the subsearch over $\mathcal{P}$, the algorithm will return projective (i.e. extremal) POVMs: our generalisation to positive matrices can thus always be thought of as a purely mathematical trick, since the measurements will always correspond to a projective measurement (and it is very important that the measurement has only 3 outcomes [34]). Importantly, no such property guarantees that the search over $\mathcal{D}$ will return pure states. The results so far correspond to a pure reference state, which caused the solution states ρ_i to also all be pure. However, one can in fact include the reference state into the search space $\mathcal{D} \rightarrow \{\rho_c, \{\rho_i\}_i\}$, thereby leaving the purity as a free parameter to be optimised over. Whilst in principle this could only lead to better results, in practice it can mean the algorithm is more likely to get stuck in local optima. However, in the appendix I present *mixed* states and measurements which lead to a $n = d = 3$ bound of $k_0 \leq 1.2018/1.5003 \approx 0.8011$, a marked improvement over known bounds with only pure states.

5.1. Noise

Inspecting Figure 2a, it seems that increasing n is an easy way to improve an experiment – but as long as the errors $\epsilon_{ijk} \neq 0$ (which captures the realistic situation where the QT predictions are not precisely reflected in the experiment) they will accumulate and spoil the trend. Let $\bar{\epsilon} = 2\sum_{i>j}[\epsilon_{ij1} + \epsilon_{ij2} + \epsilon_{ij3}]/(3n^2 - 3n)$ be the average error. Note that the objective function in (11) is monotonically increasing in $\bar{\epsilon}$. Note also that my algorithm can adjust the tradeoff between the numerator and denominator to achieve states which are in general less optimal for the noise-free case, but more robust to error than the noise-free optima. For numerical results relating to $\bar{\epsilon} = 0.15\%$, see Figure 2b.

Let a bound on k_0 be given by the ratio A/B, where A is the (noiseless) numerator of (11), and B the denominator. Then, the maximum tolerable error is $\bar{\epsilon}_{\mathrm{MAX}} = 2(B - A)/(3n^2 - 3n)$. Surprisingly, the mixed state bound I found by including the reference state in the search is extremely robust to noise, tolerating $\bar{\epsilon}_{\mathrm{MAX}} \approx 0.033$, whilst the previous best $n = d = 3$ (pure state) solution [28] could tolerate $\bar{\epsilon}_{\mathrm{MAX}} = 0.0006$. This constitutes, a ~ 50-fold improvement in robustness for the same experimental resources, contradicting the widely held feeling that extremely high precision experiments are necessary to show the reality of the quantum state.

5.2. Lower bounds

A loose lower bound on the global optimum of (11) (i.e. the best possible BCLM experiment) can be obtained through convex relaxation. Replace the bilinear terms in the numerator with zero: then the problem is convex and actually solvable analytically by putting $\rho_i = \rho_c \quad \forall i$. The BCLM argument therefore cannot hope to find a bound on

k any lower than $(1 + \frac{1}{2}n(n-1)3\bar{\epsilon})/n$. Tighter convex relaxations would be very useful if found.

6. Conclusions and directions

The technique for finding BCLM experiments presented here is very flexible. Besides the results showcased above – the improved zero-error bounds and improved finite-error bounds for low n and d – there are additional applications of the method. An experimentalist, armed with an estimate of the typical precision available in her laboratory setup, can use my algorithm with this quantity as input – and therefore extract the optimum n and appropriate sets $\mathcal{D}, \mathcal{P}$.

It is at first quite surprising that non-extremal states can be a better choice in arguing for the reality of the quantum state. It is counter intuitive because i) for many quantum information processing applications, mixed states will perform strictly worse than pure states, and ii) since mixing states together introduces 'artificial' ignorance into the problem one would then expect the states to become more epistemic. But since we are interested in the ratio of quantum to classical overlaps, we are not interested in the absolute measure of epistemicness, but rather the upper limit on how close the epistemicness of the ontological model can come to explaining the overlap at the quantum, or operational level. Intuitively, increasing the mixedness of states can make them less distinguishable: increasing their overlap ω_Q and therefore decreasing our objective function. Fundamentally, a mixed state ρ in quantum mechanics will still have to correspond to some distribution μ_ρ in an ontological model (see Figure 1). Such preparations also exhibit the property that (except in special cases) one cannot distinguish them with a single shot measurement. The epistemic interpretation is therefore just as compelling for mixed states as for pure states. One difference is that even within quantum mechanics there is an 'ignorance' interpretation of a mixed state as a (non-unique) convex combination of pure states. But since our goal is to show that there is an ever widening explanatory gap that the epistemic interpretation fails to breach, this is by no means a get out clause to the argument. Needless to say that the often-cited belief that 'there are no pure states in the laboratory' should cause ρ-epistemicism to be taken just as seriously, nay, *more seriously* than ψ-epistemicism.

Each of the mathematical elements of an ontological model (ie. the preparations, transformations and measurements) should in-principle carry labels that allow for contextuality [35, 36]. This is because multiple, physically distinct preparations (transformations, or measurements) may be identified in QT, but one is generally not warranted in identifying their representation in the ontological model. Doing so amounts to making an additional (and spurious [35, 36]) assumption of non-contextuality. It is therefore paramount in any experimental test relating to ontological models (such as those proposed by PBR and BCLM) that the very same physical procedure be used at any point where a definite preparation or measurement is repeatedly called for. The temptation to use distinct procedures which would be equivalent in QT is perhaps

greater for objects with multiple convex decompositions, such as mixed states: but it should be resisted all the same, because it would introduce a contextuality loophole.

Further, it is simple to apply additional constraints which will not spoil the properties of (11). For example; the expression $||\rho_c - \rho_j||_*$ can be upper bounded, or trace$(\rho_c \rho_j)$ can be bounded from above (and/or below if choosing ρ_c to be fixed, and not part of $\mathcal{D}$). These constraints might help provide further theoretical insight into the nature of ontological models for QT. Similar constraints might help experimentalists search for those certain states and measurements that are easy to prepare with high fidelity. It is also possible to seed the algorithm a more structured and informed initial feasible point from which to search from – rather than the random feasible points I chose. My algorithm will return a solution at least as good as its input, and can therefore be used to 'polish' any solution found by other means.

It is possible in future that deterministic global optimization techniques [37] can be applied to this problem, and provide a certificate of *global* optimality (rather than just partial optimality) for $\mathcal{D}$ and $\mathcal{P}$. Such a certificate would be a akin to a "Tsirelson's bound" [38] – and would provide some much-desired certainty in the search for optimal BCLM experiments.

Acknowledgments

I would like to thank Owen Maroney and Andrew Briggs for helpful discussions, and Ronnie Hermens and Cyril Branciard for feedback on an earlier version of this manuscript. This publication was made possible through the support of a grant from Templeton Religion Trust. The opinions expressed in this publication are those of the author and do not necessarily reflect the views of Templeton Religion Trust. This work was supported by the Royal Commission for the Exhibition of 1851.

Appendix A. Dinkelbach's technique

Dinkelbach's technique is an iterative method of solving a convex-concave fractional problem: here I follow Jeflea's description [32] – the interested reader should consult this reference for more details. Consider the problem

$$\text{minimize} \qquad \frac{a(x)}{b(x)} : x \in X, \tag{A.1}$$

where $a(x)$ and $b(x)$ are (respectively) convex and positive, concave functions and X is a compact convex set. Dinkelbach developed the following algorithm which he proved to converge on the globally optimal solution to this problem [39]:

(i) Initialise $x_1 \in X$. Let $k = 1$.

(ii) Let $\theta_k = a(x_k)/b(x_k)$. Find the solution x_{k+1} to $\min_X \{a(x) - \theta_k b(x)\}$.

(iii) If $a(x_{k+1}) - \theta_k b(x_{k+1}) = 0$, stop and x_k is optimal. Otherwise $k = k + 1$ and go to previous step.

The subproblem of my alternating search approach that involves fixing the measurements and searching for states is of the form (A.1): it may therefore be solved with Dinkelbach's method.

Appendix B. Mixed states solution

With pure states in the $n = d = 3$ case, my algorithm never returned a bound better than $k_o \leq 0.9964\ldots$ after many trials. One therefore has the suspicion that this is in fact the global optimum of the problem when the reference state is set to a pure state. Inserting the reference state into the search space removes unnecessary constraints, for example that it must have a certain purity. The algorithm found:

$$\rho_c = \begin{pmatrix} 0.36238 & -0.05820 - 0.19604i & 0.04384 - 0.07843i \\ -0.05820 + 0.19604i & 0.19657 & -0.09383 - 0.09519i \\ 0.04384 + 0.07843i & -0.09383 + 0.09519i & 0.44104 \end{pmatrix}$$

$$\rho_1 = \begin{pmatrix} 0.46135 & 0.05756 - 0.01170i & -0.08267 - 0.03289i \\ 0.05756 + 0.01170i & 0.43220 & 0.17263 + 0.06280i \\ -0.08267 + 0.03289i & 0.17263 - 0.06280i & 0.10645 \end{pmatrix}$$

$$\rho_2 = \begin{pmatrix} 0.31986 & -0.04357 + 0.08822i & 0.07957 + 0.21626i \\ -0.04357 - 0.08822i & 0.42252 & -0.07043 + 0.09595i \\ 0.07957 - 0.21626i & -0.07043 - 0.09595i & 0.25762 \end{pmatrix}$$

$$\rho_3 = \begin{pmatrix} 0.16537 & 0.09443 + 0.20002i & -0.04010 - 0.04801i \\ 0.09443 - 0.20002i & 0.34712 & 0.06583 - 0.02614i \\ -0.04010 + 0.04801i & 0.06583 + 0.02614i & 0.48752 \end{pmatrix}$$

$$U_{12} = \begin{pmatrix} -0.21308 - 0.34918i & -0.73262 - 0.03220i & 0.28060 - 0.46494i \\ 0.41066 + 0.24840i & -0.04936 - 0.42176i & -0.50882 - 0.57485i \\ -0.72704 - 0.27156i & 0.23588 - 0.47568i & -0.34016 - 0.00688i \end{pmatrix}$$

$$U_{13} = \begin{pmatrix} -0.09331 - 0.10812i & -0.82454 - 0.21487i & 0.20553 - 0.45970i \\ -0.23642 + 0.21660i & 0.25851 - 0.45038i & -0.59552 - 0.52237i \\ 0.17509 - 0.91986i & 0.03121 - 0.05766i & -0.34483 - 0.00239i \end{pmatrix}$$

$$U_{23} = \begin{pmatrix} -0.45263 + 0.11154i & -0.69730 + 0.14663i & 0.45256 - 0.26488i \\ -0.27154 - 0.06431i & 0.21650 - 0.52185i & -0.19325 - 0.75206i \\ 0.02490 - 0.83917i & -0.20971 + 0.35931i & -0.30627 - 0.16823i \end{pmatrix}$$

which gives $k_0 \leq 1.2018/1.5003 \approx 0.8011$. Note that here a measurement is represented by a unitary matrix U_{ij}, where each column is a pure state $|ijk\rangle$ corresponding to the projector $|ijk\rangle\langle ijk| = E_{ijk}$.

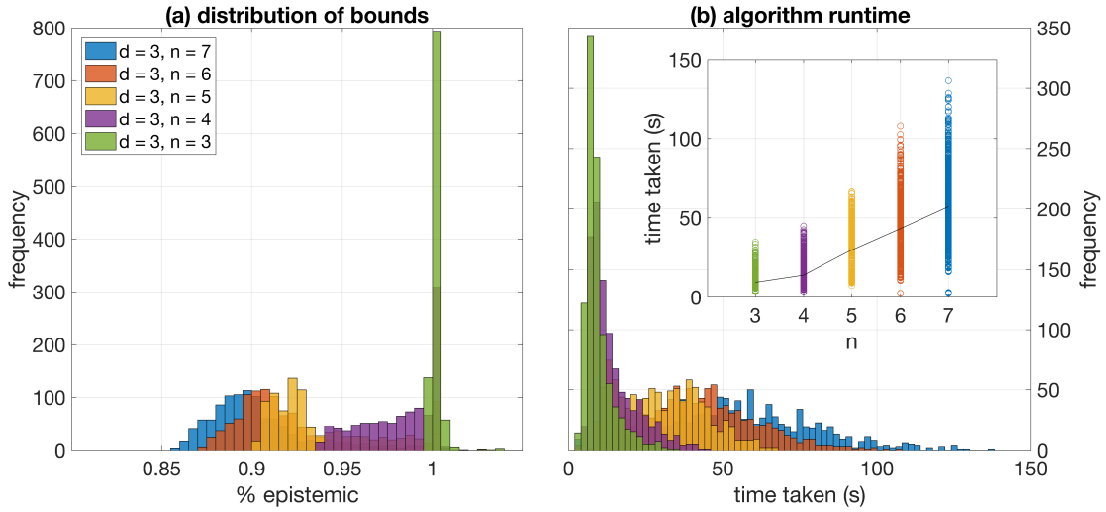


Figure C1. Histograms of the distribution of (a) output objective function value and (b) running time for 1000 trials on a single CPU implementation of the algorithm. When the algorithm is seeded randomly, it takes a variable amount of time to finish, and settles on distinct (generally non-globally optimal) solutions. The inset of (b) shows the same data with the black line tracing the mean value, which is linear for $d \geq 4$.

Appendix C. Performance of the algorithm

To gain an idea of the runtime of the algorithm, and the topology of the optimisation landscape, observe Figure C1. For the purposes of making this figure, the algorithm was run in $d = 3$, a total of 1000 times for each value of $n = 3, 4, 5, 6, 7$ on a 3.6 GHz Intel Core i7-4790 running openSUSE linux, Matlab R2016a and CVX. The results suggest that the mean runtime scales linearly with n, although generally the algorithm will find one of many distinct local optima depending on the initial random seed.

References

[1] N. Harrigan and R. Spekkens, Foundations of Physics **40**, 125 (2010).
[2] R. W. Spekkens, Phys. Rev. A **75**, 032110 (2007).
[3] M. Leifer, Quanta **3** (2014).
[4] M. Ringbauer *et al.*, Nat Phys **11**, 249 (2015).
[5] J. Bell, Reviews of Modern Physics **38**, 447 (1966).
[6] J. S. Bell, Physics **1**, 195 (1964).
[7] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt, Phys. Rev. Lett. **23**, 880 (1969).
[8] B. Hensen *et al.*, Nature **526**, 682 (2015).
[9] M. Giustina *et al.*, Phys. Rev. Lett. **115**, 250401 (2015).
[10] L. K. Shalm *et al.*, Phys. Rev. Lett. **115**, 250402 (2015).
[11] M. F. Pusey, J. Barrett, and T. Rudolph, Nat Phys **8**, 475 (2012).
[12] S. Kochen and E. Specker, Journal of Mathematics and Mechanics **17**, 59 (1967).
[13] P. G. Lewis, D. Jennings, J. Barrett, and T. Rudolph, Phys. Rev. Lett. **109**, 150404 (2012).
[14] S. Aaronson, A. Bouland, L. Chua, and G. Lowther, Phys. Rev. A **88**, 032111 (2013).

[15] D. Nigg *et al.*, New Journal of Physics **18**, 013007 (2016).
[16] M. K. Patra, S. Pironio, and S. Massar, Phys. Rev. Lett. **111**, 090402 (2013).
[17] R. Colbeck and R. Renner, Phys. Rev. Lett. **108**, 150402 (2012).
[18] L. Hardy, International Journal of Modern Physics B **27**, 1345012 (2013).
[19] O. J. E. Maroney, (2012), 1207.6906v2.
[20] J. Barrett, E. G. Cavalcanti, R. Lal, and O. J. E. Maroney, Phys. Rev. Lett. **112**, 250403 (2014).
[21] C. A. Fuchs, *Distinguishability and Accessible Information in Quantum Theory*, PhD thesis, University of New Mexico, 1996, quant-ph/9601020v1.
[22] M. Nielsen and I. Chuang, *Quantum Computation and Quantum Information (Cambridge Series on Information and the Natural Sciences)*, 1st ed. (Cambridge University Press, 2004).
[23] N. Harrigan and T. Rudolph, (2007), 0709.4266v1.
[24] L. Ballentine, (2014), 1402.5689v1.
[25] S. Bandyopadhyay, R. Jain, J. Oppenheim, and C. Perry, Phys. Rev. A **89**, 022336 (2014).
[26] C. M. Caves, C. A. Fuchs, and R. Schack, Phys. Rev. A **66**, 062111 (2002).
[27] M. S. Leifer, Phys. Rev. Lett. **112**, 160404 (2014).
[28] C. Branciard, Phys. Rev. Lett. **113**, 020409 (2014).
[29] S. R. Hedemann, (2013), 1303.5904v1.
[30] R. Enkhbat, Y. Bazarsad, and J. Enkhbayar, International Journal of Pure and Applied Mathematics **73**, 93 (2011).
[31] J. Gorski, F. Pfeuffer, and K. Klamroth, Mathematical Methods of Operations Research **66**, 373 (2007).
[32] A. Jeflea, An. St. Univ. Ovidius Constanta **11**, 87 (2003).
[33] M. Grant and S. Boyd, CVX: Matlab software for disciplined convex programming, version 2.1, 2014.
[34] A. Dutta, M. Pawłowski, and M. Żukowski, Phys. Rev. A **91**, 042125 (2015).
[35] R. W. Spekkens, Phys. Rev. A **71**, 052108 (2005).
[36] M. S. Leifer and O. J. E. Maroney, Phys. Rev. Lett. **110**, 120401 (2013).
[37] C. Floudas, *Deterministic Global Optimization: Theory, Methods and Applications*Nonconvex Optimization and Its Applications (Springer US, 2013).
[38] B. S. Cirel'son, Letters in Mathematical Physics **4**, 93 (1980).
[39] W. Dinkelbach, Management Science **13**, 492 (1967).